Tamino

Backup Guide

Version 10.15

October 2022

WEBMETHODS

This document applies to Tamino Version 10.15 and all subsequent releases.

Specifications contained herein are subject to change and these changes will be reported in subsequent release notes or new editions.

Document ID: INS-BACKUP-1015-20230910

Table of Contents

Backup Guide

The Tamino Backup Guide gives general information about backup strategies in Tamino and details about how to implement them. It is intended for Tamino database administrators with a good knowledge of Tamino administration tasks. You can make Tamino (internal) backups to disk or tape. Backing up the database to an external backup system is also available. Please consult you product support if this is considered.

This Guide covers the following topics:

1 About this Documentation

Document Conventions

Convention	Description
Bold	Identifies elements on a screen.
`Monospace font`	Identifies service names and locations in the format *`folder.subfolder.service`*, APIs, Java classes, methods, properties.
`Italic`	Identifies: Variables for which you must supply values specific to your own situation or environment. New terms the first time they occur in the text. References to other documentation sources.
`Monospace font`	Identifies: Text you must type in. Messages displayed by the system. Program code.
{ }	Indicates a set of choices from which you must choose one. Type only the information inside the curly braces. Do not type the { } symbols.
|	Separates two mutually exclusive choices in a syntax line. Type one of these choices. Do not type the | symbol.
[ ]	Indicates one or more options. Type only the information inside the square brackets. Do not type the [ ] symbols.
...	Indicates that you can type multiple options of the same type. Type only the information. Do not type the ellipsis (...).

Online Information and Support

Product Documentation

You can find the product documentation on our documentation website at **https://documentation.softwareag.com**.

In addition, you can also access the cloud product documentation via **https://www.softwareag.cloud**. Navigate to the desired product and then, depending on your solution, go to “Developer Center”, “User Center” or “Documentation”.

Product Training

You can find helpful product training material on our Learning Portal at **https://knowledge.softwareag.com**.

Tech Community

You can collaborate with Software AG experts on our Tech Community website at **https://tech-community.softwareag.com**. From here you can, for example:

- Browse through our vast knowledge base.
- Ask questions and find answers in our discussion forums.
- Get the latest Software AG news and announcements.
- Explore our communities.
- Go to our public GitHub and Docker repositories at **https://github.com/softwareag** and **https://hub.docker.com/publishers/softwareag** and discover additional Software AG resources.

Product Support

Support for Software AG products is provided to licensed customers via our Empower Portal at **https://empower.softwareag.com**. Many services on this portal require that you have an account. If you do not yet have one, you can request it at **https://empower.softwareag.com/register**. Once you have an account, you can, for example:

- Download products, updates and fixes.
- Search the Knowledge Center for technical information and tips.
- Subscribe to early warnings and critical alerts.
- Open and update support incidents.
- Add product feature requests.

Data Protection

Software AG products provide functionality with respect to processing of personal data according to the EU General Data Protection Regulation (GDPR). Where applicable, appropriate steps are documented in the respective administration documentation.

2 General Backup Strategies

Computers in general are very reliable. You may run your system for months or even years without experiencing any problems that cause you to lose information on your system. But businesses are more and more dependent on computers and the information that is stored in them. The information that is in your computer may not be available anywhere else. So every system needs to back up and restore some or all of its data. There are numerous backup strategies a company can use. In the following, you will find a short introduction to the concepts of backup, restore and recover for databases in general and Tamino in particular. The following topics are covered:

Concepts

In the following, you will find explanations of general notions and terms with regard to backing up databases. Most of them are available in Tamino, unless mentioned otherwise.

Backup

A database is saved to one or more output devices. Note that the term “backup” is used both for the process of saving the data and for the resulting data sets. Making a backup should be possible online, parallel to normal database update activities, and save a transaction-consistent state of the database. Backups should be done at a time when there is a low data load. Several backup concepts are conceivable:

Online backup: A backup during a normal update database session.

Offline backup: A backup when database updates are disabled (the server is down, or in stand-by mode, or in read-only mode).

Complete backup: All data of the whole database or the logical or physical subset of the database is saved.

Incremental backup: Only the data which has been changed since the previous backup is saved. A recovery is only possible if a previous full backup is available, as well as all following incremental backups. Incremental backup and recover operations are fast and efficient.

Full backup: The complete database is saved.

Partial backup: Only a logical or physical subset of the database is saved.

Note: Please note that partial backups are not available in Tamino.

Restore

A restore *recreates* the content of the database at backup time from the backup devices.

Full restore: The complete database is restored. A full restore is only possible after a full backup.

Partial restore: Only a logical or physical subset of a database is restored. A partial restore is possible after a full or partial backup (but not available in Tamino).

Recover from Database Logging

A database backup alone allows you only to restore one state of the database as it was at backup time. After a data failure, however, it may not be sufficient to recreate that state of a database, but the state of the database just before the failure occurred. For this reason, all update operations are logged on log spaces.

After the database has been restored, the log spaces are read and the logged update operations are repeated, so that the database is returned to the state that was valid at the time when the last log entry had been created. This process is the recover process.

Full recover: The complete database is recovered. A full recover is only possible after a full restore.

Partial recover: Only a logical subset of the database is recovered. A partial recover is possible after a full or partial restore (but not available in Tamino).

Non-Parallel Backup and Restore/Recover

Normally, a backup is performed in non-parallel mode: The data blocks are written in one stream to the backup devices, or read in one stream from the backup devices.

Parallel Backup and Restore/Recover

A parallel backup writes in parallel to more than one backup device; a parallel restore reads in parallel from more than one backup device. This increases the speed of the backup/restore process, if the backup device is much slower than the disks where the database is stored.

Replication

Conceptually different from a restore process, a database can be *created* from backup. In this case a new (duplicate) database is created with the content of the database at backup time.

Changes to the database that occur after the creation from backup can then be replicated in a replication database. Unlike a conventional backup that is restored from tape or CD, the replicated database is available to applications as soon as they can be pointed to it. For further information, see the Replication Guide.

Requirements for Backup, Logging, and Restore/Recover

After having mentioned the basic notions of backing up, the question arises why we need backup, restore and recover at all. The simple answer to that is that you want to be able to recreate a previous state of a database after an error has occurred. The reasons for errors are manifold and are dealt with in detail in the next section **Recovery from Data Failures**. Let us first consider a few requirements for being able to recreate a previous state of a database.

Database Synchronization

When a backup is performed online, it is important to be able to create a consistent state of the database after the corresponding restore. To achieve this, in Tamino a database synchronization is performed at the end of the backup: New transactions are postponed until all open transactions are finished. When all updated blocks have been written to disk, the database is in a consistent state. When this state of all database blocks is contained in the backup, it is possible to restore this (consistent) state of the database.

The restore operation recreates this consistent state of the database. Note that the restored database must be logically identical, but may be physically different. For example, the restore operation can defragment the data.

When the database server is active, it logs all update operations in the database log. After you have restored the database, the recover operation reads the logs and repeats all update operations which have been performed until the required timestamp or until the end of the logs.

Other Requirements

- If you want to be able to perform a recover after a disk error, it is necessary that the database logs and the backups are not on the same disk. This is not necessary if you only want to be able to recover from a software failure, because you have a hardware solution which guarantees that the database is not destroyed because of a hardware failure.
- Disaster Recovery (based on restore/recover) requires that the backup and log spaces be physically copied to a new computer center. For example, if the current log space is copied only after it has been closed, you are not able to reapply the changes that occurred during the current server session.

An alternative to performing a backup is just copying the database spaces to another place. But this has some disadvantages: First, it is only allowed if no update session of the database server is active. Otherwise the saved database spaces are inconsistent. Second, Tamino does not know of these “backups”. This means that old log spaces are not deleted and not released by Tamino. In addition, log spaces cannot be applied after a restore. For this reason, it is not recommended to copy the database spaces to another place instead of performing a normal Tamino backup.

Recovery from Data Failures

One of the most important tasks a database administrator has to accomplish is to define for each database how to handle data failures. There basically are four different kinds of data failure which can occur:

- Hardware Errors
- Software Errors
- Disaster
- Second Failure

Hardware Errors

A typical hardware error which may destroy a database is a disk failure. In this case, a Tamino restore/recover is a good possibility to handle the situation (see **Internal Backup and Restore in Tamino**). If you want to be able to perform a recover after a disk error, it is necessary that the database logs and backups are NOT on the same disk. Normally, a disk error is noticed as soon as it occurs. Hardware errors, that are not immediately recognized, are more problematic, for example if a disk read operation does not display an error, but returns wrong data. This situation is similar to software errors (see next section **Software Errors**). Another solution for handling disk errors is saving backups on a tape.

There may be other hardware errors which do not require a restore/recover, but for example a new database start to be performed after repairing the hardware. Note that there are also other concepts of handling disk errors, for example RAID 5 or cluster solutions:

- RAID 5 or disk mirroring: The data is stored redundantly on the disks. If a disk is corrupted, the data is automatically read from other disks. The computer operator must only replace the corrupted disk. The data is automatically recreated on the new disk.
- Replication: After a hardware error has occurred, a replication of the database becomes the master database. It must be made available with the name of the original database. The advantage of this solution is that the database is available without time losses required for a restore/recover process. On the other hand, some transactions may be lost in this process.

The following table compares various possibilities available to recover from a hardware error:

Solution	Special Hardware Requirements	Recovery Time	Loss of Data	Remarks
Internal Restore/Recover in Tamino	None	Long	No	-

Solution	Special Hardware Requirements	Recovery Time	Loss of Data	Remarks
RAID 5 or disk mirroring	There are hardware or software based solutions, where the operating system manages the disks	None (the user does not notice that there is a disk failure)	No	If the system is not based on physically separated storage devices, an additional recovery solution should be provided in case the whole storage system fails.
Tamino Replication	None	Short, but in contrast to high availability, the replication database must be made available as the master database manually.	Yes; because the replication is done asynchronously, the last transactions may be lost.	This solution allows also recovery from other hardware errors, for example CPU failures.

Software Errors

Contrary to recovery from hardware errors, an automatic recovery from software and handling errors is not possible. For the system, a software error is like a normal update operation. The database administrator has several possibilities for handling the problem:

- Perform a restore/recover to a state before the error occurred.
- Tamino software error: In some cases it might help to restore a backup created before the erroneous Tamino version was installed, and to recover all logs with a Tamino server in which the problem had been fixed. (Note that it is possible to restore a backup from a former Tamino version, even if that Tamino version is not installed).
- Try to repair the error, for example by updating the corrupted data or unloading the data that is not corrupt, deleting the corrupted data and reloading the correct data.

The solution depends very much on the individual error situation. Nevertheless, it is useful to perform regular backups, so that backup and restore/recover is a feasible possibility in each situation.

Disaster

It may happen that not only part of the hardware is erroneous, but that the whole hardware system is destroyed, even the complete computer center. In this case, it is necessary to make the data available on another computer, in a different place. This scenario is called disaster recovery. Concepts of disaster recovery are not necessarily based on backup and restore mechanisms. You can also use replications or cluster solutions with physically distributed storage devices. In any case all data required for the disaster recovery must be saved at a remote location. The following table shows the various possibilities you have for disaster recovery in Tamino:

Solution for Disaster Recovery	Special Hardware Requirements	Required Recovery Time	Remarks
Tamino Restore/Recover	None	Long	The updates of the current logs are lost if the log spaces are only copied after they have been finished.
Disk mirroring on remote location	Yes, but possibly there are also software-based solutions available	Short, the server needs only to be started on the target machine.	The same precautions as for a cluster solution are necessary.
Tamino Replication	No	Short, but the replication database must be manually made available as the destroyed master database.	Updates may be lost!

For more information about disaster recovery in Tamino, see the section ***Disaster Recovery*** in this guide and the documentation about *High Availability*.

Second Failure

In addition to a single hardware error, the database administrator must be aware of the fact that there is a small risk of a second failure. Standard backup solutions guarantee recovery only if not more than one disk crashes. However, if for example you perform an internal backup, a disk containing a database space has crashed, and the backup is not readable, the complete data is lost.

Depending of the kind of recovery, the following strategies can be provided in case of a second failure:

- If you have a separate solution for disaster recovery, you can use this solution for a second failure. But be aware of the fact that if the only solution for recovery from hardware failures is the disaster recovery solution, this may not be sufficient.
- If you are performing internal backups, you can use a previous backup. In this case, it is important that the different backups are stored on different physical devices. If you want to be able to also perform a recover process in the case of a second failure, you must also save the log spaces to different physical devices. If you restore an older backup, recovery will take longer than usual, because also the logs created between this backup and the backup which is no longer readable must be applied. You can avoid this by copying the backups to another physical location.
- Logs should be copied if you do not want to lose updates because logs are no longer readable.
- If you use RAID technology, you can additionally perform backups and copy the backups to another device. You will usually reduce the frequency of the backups, for example once a week instead of once a day.

Time Considerations

In the day-to-day administration environment, there normally is a requirement stating that the database should be up after a failure within a given amount of time, for example within two hours. This means that the database administrator must estimate the time necessary for a restore/recover process. Total recovery time is the result of summing up the restore time and the recover time. Use the figures given in the following rule-of-thumb example, to calculate an estimate of the restore/recover time.

Assume the estimated update time is half an hour and the estimated recover time for the updates of one day is a quarter of an hour, and that you have 5 working days with update activity. Assume the restore/recover time should be no more than 2 hours, after which you should do a weekly backup. If the failure occurs very shortly after the backup, the restore/recover time would be about half an hour. If the failure occurs after one week, the restore time would be about 0.5 h + 5 * 0.25 h = 1.75 h. In this case, it is recommended to perform a weekly backup. There may be 20% more update activities than usual, and the restore/recover time is still not more than 2 hours.

Note the following rules-of-thumb:

- The restore time is proportional to the backup time. Compare the backup and restore time and use the resulting factor: When the database grows, you can estimate the restore time by multiplying the current backup time by the factor.
- The recover time is proportional to the size of the log files, unless mass loads or index creation operations must be recovered. Also, the recover time for a given amount of log files may vary, depending on the number of update operations.

3 Internal Backup and Restore in Tamino

The following topics describe how to internally back up and restore your Tamino database in case of any failure that can occur on your system: disk failure, program failure, or human error. When you perform an *internal* backup in Tamino, all data in the database is written block by block. The following topics are covered:

Backup

General

An internal Tamino backup can be a full backup of the whole database, or it can be an incremental backup that records the changes since the last full or incremental backup. When developing a backup strategy, apply this principle: the more often your database changes, the more often you should save it. If, for example, your database is very active, you should make daily backup copies. The backup is done online, which means that the server is running, so you do not even have to stop the database when backing up, or notify your users. After a full or incremental backup, a new Tamino log space is started automatically. Backups can also be done in read-only or stand-by mode. This kind of backup that is done while it is not possible to update the database is called *offline backup*.

Incremental backups provide a useful extension to Tamino's backup capabilities. If you decide to include incremental backups in your backup strategy, you should keep a sensible balance between incremental backups and full backups. Regular full backups should still be the central part of your strategy, and the incremental backups can be considered as a quick and safe way of protecting against data loss in the intervals between the full backups.

Generally, internal backups are done to another disk. If you have enough space on a disk other than your working disk, you should back up to a disk as your preferred method. This saves time and effort.

Caution: Please note that when you delete a Tamino database, all backups of the database are also deleted if the option **Keep Backups** has not been specified.

How to Back Up

To create a backup of a database

- Follow the steps described in section Database Backups in the documentation of the Tamino Manager.

Backup Generations

One important aspect to consider when backing up is the number of backup generations that may co-exist (the default value is 5). If you have enough space on your storage device, a recommendation for the number of backup copies is 7, which means one backup for each day of the week. After the seventh backup, the first one will be deleted, including the log spaces and everything else belonging to the backup.

Why should you have more than one backup? The last backup is always taken as the default, because it takes the least time for the subsequent Recover step. The backup(s) before that should only be used if the most recent backup is not usable. Security increases with the number of backup copies available.

If you use incremental backups as part of your backup strategy, note that incremental backups are not counted as additional backup generations. A backup generation can be thought of as a full backup plus all of the associated incremental backups. A new backup generation begins when a new full backup is made.

Note: Note that newer backups will be removed when an old backup had been restored and the database has been modified.

Disabled Backups

If you perform a database restore without recover, using a backup that is not the most recent backup, all subsequent backups of the database are set automatically to disabled.

Disabled backups are ignored when Tamino counts the number of backup generations for a particular database.

You can use a disabled backup to create a database, using the command `Create Database from Backup`, but otherwise a disabled database cannot be used any more.

Note: There are currently no direct features for deleting a disabled backup. However, a disabled backup will be automatically deleted when there are no more older non-disabled backups.

Restore and Recover

To be prepared for certain types of disk failures, you need to be able to recover everything on your system.

The process of regenerating your data from a backup is done in two steps: First, the restore step, which is the reverse procedure of a backup. As input, you use the backup copy of the database, which automatically restores your data back into the database. The second step, the recover step, repeats database changes since that last backup by using the session logs stored in the log spaces.

A restore is always done for the whole database, not for parts of the database, e.g. for single collections. When you restore from a full backup, the database contents are returned to their exact state at the time the full backup was made. If you restore from an incremental backup, Tamino first restores the most recent full backup that was made before the incremental backup, then restores any other incremental backups that were made after the full backup but before the selected incremental backup, then finally restores the selected incremental backup.

As soon as the restore step successfully finishes, the database is automatically started in standby mode. Then the recover step is started. During the recover step, all of the completed database transactions that were made since the specified backup are recovered. The input for the recover step are the log spaces; each server session creates a log space that records all completed database transactions that are made during the session. By using the **Recover Until** option, you can restrict the recovery until a certain date and time.

If you restore from a full backup, the subsequent automatic recover step uses only the session logs, even if there are incremental backups available that were made after the full backup.

If you restore from an incremental backup, then Tamino restores first the corresponding full backup and then all incremental backups up to and including the selected incremental backup. The subsequent automatic recover step uses the session logs that were created after the selected incremental backup. Incremental backups that were made after the selected incremental backup are not used in the restore or recover phase.

Tamino processes incremental backups faster than session logs. Therefore, if you have made a full backup and one or more subsequent incremental backups, then you should select the most recent incremental backup for doing the restore instead of selecting the full backup.

The log spaces are located in the folder "<INSTALL_ROOT>/Tamino/db". During a mass loading process, the necessary information is written to so called “Utility Recovery Spaces”. These spaces are handled like log spaces during the recovery step.

An important aspect when restoring and recovering data is the consistency of the database. In Tamino, you can be assured that even if it was not possible to recover all changes, the database is consistent. Transactions are recovered in an “all-or-nothing” fashion. The log spaces carry a

timestamp so that you can recover your data up to a certain point in time, e.g. before an application error occurred.

Also, it is possible to restore a backup of a former Tamino version, whether the older Tamino version is installed or not.

How to Restore and Recover

Note: During the Restore/Recover step the database is in standby mode, which means that it is not possible to work with it. Hence before and after changing large amounts of data at one point in time, it is recommended to back up the database. Otherwise, the Restore/Recover step might take very long, and the database is not available for that period of time.

How to handle Restore and Recover Error Situations

When you do a Recover, you have the choice between three possibilities: **Recover all**, **Recover until**, and **Don't Recover**. In an error-free situation, you naturally would recover all changes made to the database since the last backup. The other two options are to be used if an application error occurs, or if the Restore/Recover step is not successful.

Restore Error Situations

Restoring data is just the reverse of backing up data. The data saved in the backup copy is put back to where it belongs. Possible error situations for restoring database backup copies are for example:

- The last backup copy was accidentally deleted, or cannot be found.
- The last backup copy is corrupt.

In all these cases, it is important to have a good backup strategy. If you have several backup copies, go back to the most recent one that is not corrupt, and do a Restore/Recover.

Recover Error Situations

For the Recover step, though, other error situations may occur:

- Some or all log spaces are missing for the Recover step.
- The server process is aborted while the Recover process is running.

If there are no or not all log spaces available:

The first case might occur due to a hardware failure or to human error. Maybe someone has deleted or renamed the log spaces accidentally, or the directory no longer exists. If at least some of the log spaces for the Restore/Recover process are available, Tamino processes that information up to a point in time just before information is missing or corrupt. In order to keep the database consistent,

transactions that were still open when the error occurred are rolled back. After having recovered the available information, the server is stopped and shut down. The database is consistent, but not as current as it was before the error occurred. Notify your users to update the database manually.

If no log spaces at all can be found, it is not possible to recover any changes. In this case, the only possibility is to check the **Don't Recover** option and to just restore the database, without applying the changes since the last backup. Again, notify your users to update changes manually.

If the server process is aborted while the Recover step is running:

In the second case mentioned above, you have started a Recover step, but it is interrupted because the server is no longer available, e.g. due to a power failure. When you restart the server, several messages are displayed, informing you about the error status of the Recover step. At a certain point in time, you will be asked if you want to continue the server startup:

```
Please reply "yes" or "no" to continue server startup
```

The default is `no`. Reply `no` if you do not expect that the Recover step will be aborted again (for example, if the reason for the abort was a power failure). Redo the Restore/Recover step.

Reply `yes` if an abort will probably reoccur with a second trial, for example if the reason for the abort was a corrupted log space. In this case, you will be informed that restore of existing backups created after the one used for the failed Restore/Recover run is not allowed any more. It is still possible to retry Restore/Recover for the same or older backups. For safety reasons, it is recommended to create a new backup immediately. Alternatively, you could also reply `no` and do a `Recover until`. Specify a point in time just before the server crash occurred, so that the corrupted log space which contains the error is not used.

To Recover until a certain point in time

- use the option **until-date-time** in the **inoadmin** commandline tool, as described in section Database Backups. Enter the date and time until which you want the Recover step to run.

Note: You can use the **Recover until** option also if you are sure that an error occurred, e.g. due to a faulty application, or if you want to delete parts of the changes that were made to a database and cannot be deleted otherwise.

Summary

A backup always applies to a whole database. If you back up your data daily and keep your backups in a safe place, you cannot lose more than a day's work. With the Restore/Recover functions you have the possibility to regenerate changes made to the database after the most recent backup. The database will always be restored to a consistent state, but in very rare cases, errors might occur during the process. If an error occurs, there are two possibilities: Either apply changes only up to a certain point in time, or skip the whole recover step and just restore data back to the latest backup copy. In both cases, notify your users to update the database manually with the changes that were lost.

4 Disaster Recovery

In Tamino, disaster recovery stands for the restoration of a Tamino database when everything is lost except for a backup of the database and recovery files. This may be the case after a total loss of the computing environment, in which backup and recovery files are still available on an archive. Unlike the normal restore and recover process, no information about the locations of the recovery files exists within the internal configuration store. Nevertheless, Tamino offers the possibility to recover from this disaster scenario if the following conditions are met:

- The log spaces since the last backup used to create a database must be complete (no log space must be missing).
- All log spaces must be accessible.
- The log spaces must not have been used for a **Recover until** or **Skip recover** process before.

  In order to meet these conditions, you should have a good archiving concept. Please contact Software AG Support, if you require support for this task.

The following two disaster recovery scenarios are described in this chapter:

Recovering from Internal Configuration Store Corruption

If the content of your internal configuration store (registry) is damaged for any reason (e.g. someone has accidentally deleted the information, or a disk error has occurred), you will not be able to access the Tamino databases. However, the file *Registry.log* keeps track of all internal configuration store entries that are relevant to Tamino and have occurred since the last installation of Tamino.

On Windows

To recover from registry corruption

1 Use the import file provided by the installation of the latest Tamino version

```
regutil delete "HKLM\Software AG\Tamino"

regutil import "<DRIVE_LETTER>:\Documents and Settings\All Users\Application
Data\Software AG\Tamino\Registry.export"
```

2 Apply the internal configuration store entries from the saved log with

```
regutil -f "<DRIVE_LETTER>:\Documents and Settings\All Users\Application
Data\Software AG\Tamino \Registry.log"
```

On UNIX

Case 1: A backup of $SAG/common/rgs/REGFILE is available

1 Restore the backup of the internal configuration store

$SAG/common/rgs/REGFILE

2 Apply the internal configuration store entries from the saved log with

```
regutil -f "$SAG/ino/Registry.log"
```

Case 2: A backup of $SAG/common/rgs/REGFILE is not available

1 Save the file "*$SAG/ino/Registry.log*"

2 Reinstall Tamino, including the internal products.

3 Apply the internal configuration store entries from the saved log with

```
regutil -f <saved Registry.log>
```

Note: Note that the backup file of the *Registry.log* must be a backup of the current Tamino installation, and not of a former version of Tamino.

Recovering from the Loss of the Host Environment

If you have lost your total computing environment due to a disaster, it is still possible to recover if a backup and all subsequent recovery files are available. However, the recovery files must not have been used for a restore with no full recover (recover with Skip Recover or Recover Until) process. The steps for disaster recovery in this case are:

To recover from the loss of the host environment

1 In a new computing environment, install Tamino.

2 Perform a **Create from Backup** with the given backup file.

Note: The backup file you specify must be a full backup file. You cannot use an incremental backup file to perform this step.

If no recovery files are to be applied or available, this is all that needs to be done. This may be the case if you are doing frequent backups.

If recovery files are available, apply them by following the next steps:

3 Add the internal configuration store entries for the recovery files.

To do so, create a script using the `regutil` utility to add the entries for the recovery files. An entry for a log space (with the extension .1LO) has the following structure:

```
HKEY_LOCAL_MACHINE
\SOFTWARE
 \Software AG
  \Tamino
   \servers
    \<database name>
     \savepoints
      \000000000000
       \log space
        \<session number>
         \<extent number>
```

The session number must be extended by leading zeroes as a 10 digit number, the extent number as a 5 digit number.

Here is an example: Two log spaces, AAF000010000000025.1L0 and AAF000010000000026.1L0 on the directory D:\archive, are to be processed.

```
create  "HKLMS\Software AG\Tamino\servers\db1\savepoints\000000000000\log space"

create      "HKLMS\Software AG\Tamino\servers\db1\savepoints\000000000000\log ↩
space\00025"

create      "HKLMS\Software AG\Tamino\servers\db1\savepoints\000000000000\log ↩
space\00025\00001"

setvaluedata "HKLMS\Software AG\Tamino\servers\db1\savepoints\000000000000
              \log space\00025\00001" type "file"

setvaluedata "HKLMS\Software AG\Tamino\servers\db1\savepoints\000000000000
             \log space\00025\00001" path "D:\archieve\AAF000010000000025.1L0"

create       "HKLMS\Software AG\Tamino\servers\db1\savepoints\000000000000
              \log space\00026"

create       "HKLMS\Software AG\Tamino\servers\db1\savepoints\000000000000
              \log space\00026\00001"

setvaluedata "HKLMS\Software AG\Tamino\servers\db1\savepoints\000000000000
              \log space\00026\00001" type "file"

setvaluedata "HKLMS\Software AG\Tamino\servers\db1\savepoints\000000000000
              \log space\00026\00001"

path         "D:\archieve\AAF000010000000026.1L0"
```

Other recovery files than log spaces (extension .1L0) only exists when the Tamino Data Loader was used without the `concurrentWrite` option. These are files with an extension of .1C0, .1C1 or .1C2. In case you want to do a disaster recovery with Tamino Data Loader recovery files, please contact Software AG Support to ask how to add registry entries.

4 Use the Restore/Recover function.

Note: Note that the server must not be started between **Create from Backup** and performing the recover function.

Considerations for Incremental Backup

The incremental backup feature is selectable for each database individually, i.e. you can activate it for one database and deactivate it for another. When you use the Tamino Manager to activate or deactivate the feature for a particular database, the change takes effect the next time the database server is restarted.

To specify that a database can be used for making incremental backups, do the following:

To specify that a database can be used for making incremental backups

1 In the Tamino Manager, select the node of the database.

2 Expand the tree of the node.

3 Expand the **Properties** node.

4 In the context menu of the **Server** node, select **Modify**.

5 Set the value of the property `incremental backup` to "yes" and choose **OK**.

The following restrictions apply for incremental backups:

- Incremental backup is not available for external backup environments.
- The parallel backup feature is not available for incremental backups.
- Incremental backups are not available for read-only databases.

Using Backup Generations

Make sure to develop a backup strategy that gives you maximum protection for your data. If you have many update operations per day, it is recommended that you also back up the database at least daily. The number of backup generations to be retained can be defined using the backup parameter **number of backup generations**. If you maintain for example seven backup generations, you have one backup for each day of the week. When you create the 8th backup, the first one will be deleted.

Note: Backup generations apply to full backups only, not to incremental backups.

Using the **inoadmin** tool the number of backup generations can be viewed via the **getproperty** command and modified via the **setproperty** command.

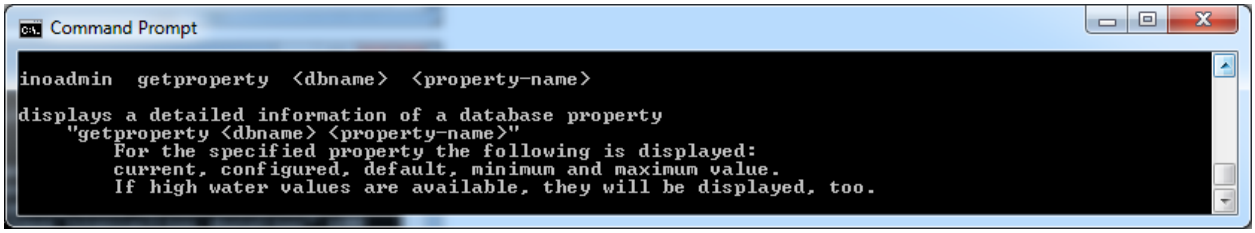

The property that rules how many backups are kept is called **number of backup generations**. The default is five.

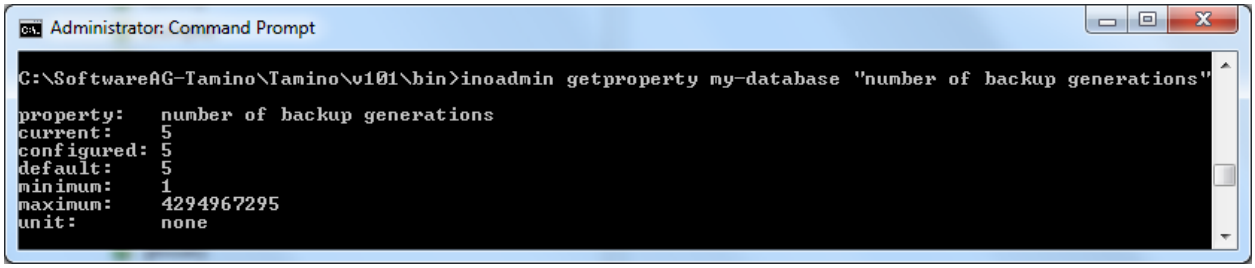

The property can be modified via the **setproperty** command. Use the **restart** option to make the property change effective at once.

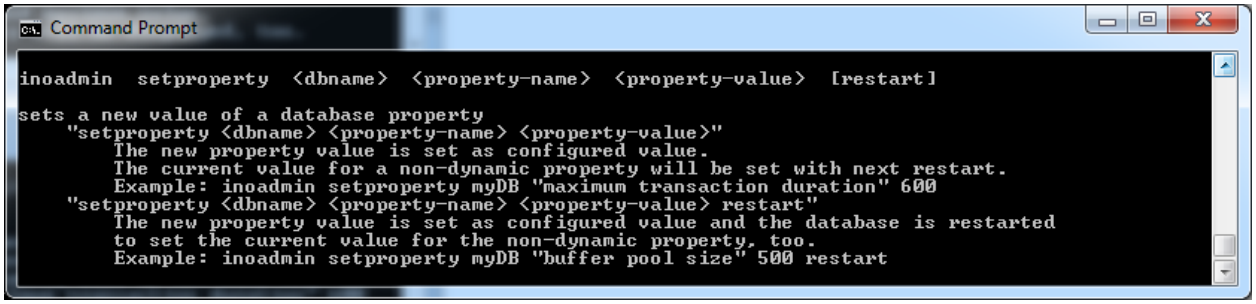

The following sets the property to seven in order to have one backup for each day of the week.

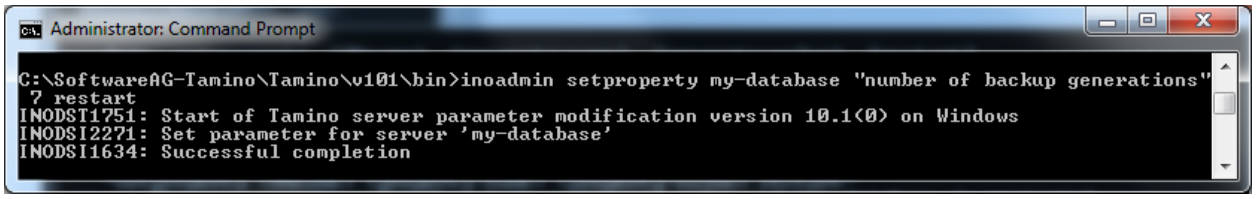

Refer also to the section Internal Backup and Restore in Tamino in the *Backup Guide* for more issues concerning backup generations.

Index